



Open NAND Flash Interface Specification: Block Abstracted NAND

BA NAND Revision 1.0
18-July-2007

Hynix Semiconductor
Intel Corporation
Micron Technology, Inc.
Phison Electronics Corp.
Sony Corporation
STMicroelectronics

This 1.0 revision of the Open NAND Flash Interface Block Abstracted NAND specification ("Final Specification") is available for download at www.onfi.org.

SPECIFICATION DISCLAIMER

THIS SPECIFICATION IS PROVIDED TO YOU "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, NON-INFRINGEMENT, OR FITNESS FOR ANY PARTICULAR PURPOSE. THE AUTHORS OF THIS SPECIFICATION DISCLAIM ALL LIABILITY, INCLUDING LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS, RELATING TO USE OR IMPLEMENTATION OF INFORMATION IN THIS SPECIFICATION. THE AUTHORS DO NOT WARRANT OR REPRESENT THAT SUCH USE WILL NOT INFRINGE SUCH RIGHTS. THE PROVISION OF THIS SPECIFICATION TO YOU DOES NOT PROVIDE YOU WITH ANY LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS.

Copyright 2005-2007, Hynix Semiconductor, Intel Corporation, Micron Technology, Inc., Phison Electronics Corp., Sony Corporation, STMicroelectronics. All rights reserved.

For more information about ONFI, refer to the ONFI Workgroup website at www.onfi.org.

All product names are trademarks, registered trademarks, or servicemarks of their respective owners.

ONFI Workgroup Technical Editor:

Amber Huffman
Intel Corporation
2111 NE 25th Ave M/S JF2-53
Hillsboro, OR 97124 USA
Tel: (503) 264-7929
Email: amber.huffman@intel.com

Table of Contents

| | | |
|---------|--------------------------------------------------|----|
| 1. | LBA Abstraction | 4 |
| 1.1. | Definitions, abbreviations, and conventions..... | 4 |
| 1.1.1. | Definitions and Abbreviations | 4 |
| 1.2. | Timing requirements | 4 |
| 1.2.1. | Unexpected Power Fail..... | 4 |
| 1.3. | Write Protect | 5 |
| 1.4. | Data and Metadata Reliability | 5 |
| 1.5. | Command Set | 5 |
| 1.6. | Reset Definition..... | 6 |
| 1.7. | Read Parameter Page modifications | 6 |
| 1.7.1. | Additions to Parameter Page Data Structure | 6 |
| 1.8. | LBA Read Definition..... | 8 |
| 1.9. | LBA Write Definition | 11 |
| 1.10. | LBA Read Metadata Definition | 14 |
| 1.11. | LBA Deallocate Definition..... | 16 |
| 1.12. | LBA Flush Definition | 17 |
| 1.13. | LBA Abort Definition | 18 |
| 1.14. | Read Status Definition..... | 19 |
| 1.15. | Read Unique ID Definition..... | 19 |
| 1.16. | Status Field Definition..... | 20 |
| 1.17. | Feature Parameter Definitions | 21 |
| 1.17.1. | BA NAND: Error Information | 21 |
| 1.17.2. | BA NAND: Configuration..... | 22 |

1. LBA Abstraction

The device may support Block Abstracted access mode. In this mode, all data is accessed by logical block address (LBA) in granularity of a sector. The Block Abstracted device handles ECC and NAND management functions. The host may issue reads and writes for multiple LBAs, with flow control capabilities that the device may employ in place. Block Abstracted access mode devices support 8-bit data buses only.

The logical block address is in units of 512 bytes. The LBA is always 40 bits, resulting in five address cycles for transmission of the LBA.

This specification is a supplement to and should be used in conjunction with the ONFI 1.0 specification. Packages and pinouts, timing diagrams and parameters, and commands required but not defined in this document are as defined in the ONFI 1.0 specification.

1.1. Definitions, abbreviations, and conventions

1.1.1. Definitions and Abbreviations

The terminology used in this specification is intended to be self-sufficient and does not rely on overloaded meanings defined in other specifications. Terms with specific meaning not directly clear from the context are clarified in the following sections.

1.1.1.1. logical block address (LBA)

The logical block address (LBA) is an address in fixed units of 512 bytes. The LBA is used in commands to identify the sector(s) to access. The LBA shall always be specified on a sector boundary.

1.1.1.2. sector

A uniquely addressable unit of data. This unit of data is always a power of two and not smaller than 512 bytes.

1.2. Timing requirements

The device shall complete its power on reset initialization, t_{PWR_ON} , within 250 milliseconds after Vcc reaches Vcc_min. Specifically, R/B# shall be set to one no later than 250 milliseconds after Vcc reaches Vcc_min.

The device shall complete a Reset command, t_{RST} , prior to any data transfers within the power on reset initialization time requirement (t_{PWR_ON}). The device shall complete a Reset command, t_{RST} , after any data transfers have been issued within the LBA Flush time requirement (t_{FL}).

1.2.1. Unexpected Power Fail

If an unexpected power failure occurs, the device may need time to recover upon the subsequent power on. An unexpected power failure is when a power loss occurs that was not preceded by an LBA Flush with the Standby parameter set to one as the preceding command.

The device reports that it is recovering from an unexpected power failure in the PFR bit of the Status register. This bit may be set to one while the device recovers during the next power on. The power fail recovery shall finish no later than one minute after R/B# is initially set to one by the device.

While the device is performing power fail recovery, the host shall not issue the following commands: LBA Read, LBA Read Continue, LBA Write, LBA Write Continue, LBA Read

Metadata, LBA Deallocate, LBA Flush, or LBA Abort. The device shall respond to and execute the following commands if received from the host during power fail recovery: Read Status, Read ID, Read Parameter Page, Read Unique ID, Get Features, Set Features, and Reset.

1.3. Write Protect

Write Protect is not used by devices that support Block Abstracted access mode. The WP# signal is ignored by a Block Abstracted device (it is treated as a don't care). Note that the host may use WP# since prior to enumeration completing the host does not know whether Block Abstracted access mode is supported.

1.4. Data and Metadata Reliability

The host does not need to perform any error correction on data returned by a device that supports Block Abstracted access mode. The device shall return only valid data and metadata. If the device cannot return valid data or metadata, the device shall fail the corresponding command by setting the FAIL bit to one in the Status register.

1.5. Command Set

Table 1 outlines the ONFI command set for a device that supports Block Abstracted access mode.

The value specified in the first cycle identifies the command to be performed. For commands that include parameters or data input, there is a second cycle value that specifies that the command shall start execution after the parameters/data are specified. For commands that do not include parameters or data input, the command shall start execution after the first cycle; the second cycle is not applicable for such commands.

| Command | O/M | 1 st Cycle | 2 nd Cycle | Acceptable while Busy |
|---------------------|-----|-----------------------|-----------------------|-----------------------|
| LBA Read | M | C0h | 30h | |
| LBA Read Continue | M | C8h | | |
| LBA Write | M | C1h | 10h | |
| LBA Write Continue | M | C2h | 10h | |
| LBA Read Metadata | O | C0h | 31h | |
| LBA Deallocate | M | C3h | 10h | |
| LBA Flush | M | C9h | | |
| LBA Abort | M | CAh | | Y |
| Read Status | M | 70h | | Y |
| Read ID | M | 90h | | |
| Read Parameter Page | M | ECh | | |
| Read Unique ID | O | EDh | | |
| Get Features | M | EEh | | |
| Set Features | M | EFh | | |
| Reset | M | FFh | | Y |

Table 1 Command set

1.6. Reset Definition

The Reset function puts the device in its default power-up state. Note that some feature settings are retained across Reset commands. The command may be executed with the device in any state. Figure 1 defines the Reset behavior and timings.

For devices that support Block Abstracted access mode, the Reset command does ensure that all data previously written to the device is placed into the Flash array prior to performing the reset of the device.




Figure 1 Reset timing diagram

1.7. Read Parameter Page modifications

Block Abstracted access mode devices are allowed up to 10 milliseconds to retrieve the parameter page and make it accessible to the host.

Devices that support Block Abstracted access mode are only required to have one parameter page. Block Abstracted devices that can guarantee correct data is returned do not need to store redundant parameter page copies.

1.7.1. Additions to Parameter Page Data Structure

Unless specified in this document, the definition in the ONFI 1.0 specification for a field in the parameter page applies. For parameters that span multiple bytes, the least significant byte of the parameter corresponds to the first byte.

| Byte | O/M | Description |
|-------------------------------------------------------------------------------------------------|-----|------------------------------------------------------------------------------------------------------------|
| Features block | | |
| 6-7 | M | Features supported 8-15 Reserved (0) 7 1 = supports Block Abstracted access mode 0-6 Reserved (0) |
| 8-9 | M | Optional commands supported 6-15 Reserved (0) 5 1 = supports Read Unique ID 0-4 Reserved (0) |
| LBA organization block (Only present when Block Abstracted Access Mode is supported) | | |
| 80-87 | M | Number of LBAs |
| 88-89 | M | Sector Size |
| 90-91 | M | Sector Multiple |
| 92 | M | Metadata Bytes Per Sector |
| 93-127 | | Reserved |
| Timing parameters block (Only present when Block Abstracted Access Mode is supported) | | |
| 133-134 | M | t_{RL} Maximum LBA read time (ms) |
| 135-136 | M | t_{WL} Maximum LBA write time (ms) |
| 137-138 | M | t_{FL} Maximum LBA flush and runtime reset time (ms) |
| 139-163 | | Reserved |

Table 2 Parameter page definitions

1.7.1.1. Byte 6-7: Features supported

This field indicates the optional features that the device supports.

Bits 0-6 are reserved and shall be cleared to zero.

Bit 7 when set to one indicates that the device supports the Block Abstracted access mode type, and supports the LBA Read and LBA Write commands. When cleared to zero, the device does not support the Block Abstracted access mode type, and instead supports the Read, Page Program, and Erase commands. The device shall only support one access mode type.

Bits 8-15 are reserved and shall be cleared to zero.

1.7.1.2. Byte 8-9: Optional commands supported

This field indicates the optional commands that the device supports.

Bits 0-4 are reserved and shall be cleared to zero.

Bit 5 when set to one indicates that the target supports the Read Unique ID command. If bit 5 is cleared to zero, the host shall not issue the Read Unique ID command to the target.

Bits 6-15 are reserved and shall be cleared to zero.

1.7.1.3. Byte 80-87: Number of LBAs

This field defines the total number of logical blocks (LBAs) present in the device. This value shall be the total user storage capacity of the device in bytes divided by 512.

1.7.1.4. Byte 88-89: Sector Size

This field indicates the sector size for logical blocks, when the Block Abstracted access mode is supported. The value is reported in terms of a power of 2. For instance, a reported value of 14 corresponds to a sector size of 2^{14} or 16384 bytes. A value smaller than 9 (i.e. 512 bytes) is not supported.

1.7.1.5. Byte 90-91: Sector Multiple

This field indicates the Sector Multiple value. When doing LBA Read or LBA Write commands that have sector counts larger than one, the transfers shall always be broken into transfer sizes that are smaller or equal to Sector Multiple sectors by the host. Between each Sector Multiple sectors, an LBA Read Continue or LBA Write Continue command is issued to the device which allows the device to perform flow control. The Sector Multiple may also indicate the preferred write size transfer length (or multiple thereof) to provide high write performance.

1.7.1.6. Byte 92: Metadata Bytes Per Sector

This field indicates the number of bytes available for metadata per sector. The metadata is read and written atomically with the corresponding data sector. If the value in this field is 0h, then the Block Abstracted device does not include any metadata space per data sector and metadata shall not be used. If the value in this field is non-zero, then the device shall support the LBA Read Metadata command.

1.7.1.7. Byte 133-134: Maximum LBA read time

This field indicates the maximum LBA read time (tRL) in milliseconds.

1.7.1.8. Byte 135-136: Maximum LBA write time

This field indicates the maximum LBA write time (tWL) in milliseconds.

1.7.1.9. Byte 137-138: Maximum LBA flush time

This field indicates the maximum LBA flush time (tFL) in milliseconds. This value is also the maximum runtime reset (tRST) time in milliseconds.

1.8. LBA Read Definition

The LBA Read function reads sectors of data identified by an LBA and a sector count. The logical block is made available to be read from the data latch. If metadata use is enabled, then the metadata bytes for an addressed sector are returned immediately following the data for the sector. Thus, metadata if used is interspersed between the data sectors returned. Figure 2 defines the LBA Read behavior and timings for the transfer of the first data block.

Read Status shall be used to check the status of LBA Read (Continue). The host shall ensure that FAIL is cleared to zero in the Status register which indicates the data read was successful prior to transferring data. After completion of the Read Status command, the C0h command shall be issued by the host to continue with the data output flow for the LBA Read (Continue) command.

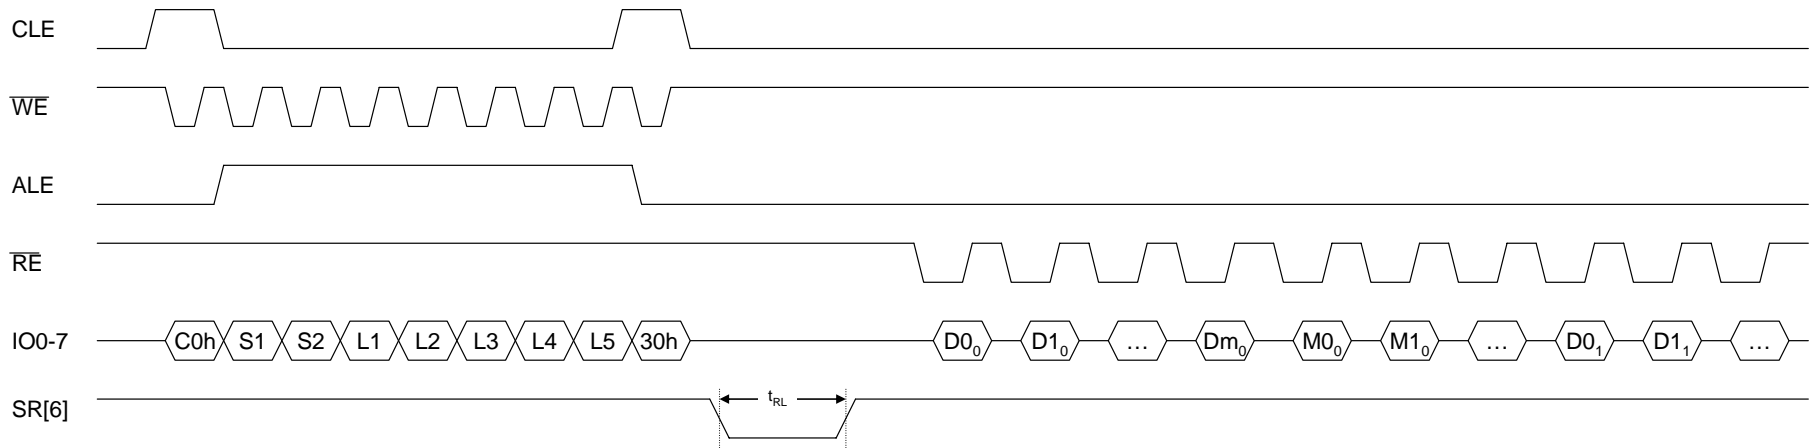


Figure 2 LBA Read timing

L1-L5 LBA of the sector(s) to retrieve. L1 is the least significant byte.

S1-S2 Number of sectors to return to the host. S1 is the least significant byte.

Dn₀ Data bytes read from the first addressed sector.

Dn₁ Data bytes read from the second addressed sector.

Mn₀ Metadata bytes read for the first addressed sector.

When many sectors are being transferred, the host shall issue the LBA Read Continue (C8h) command every time Sector Multiple sectors have been transferred. All data for the previous Sector Multiple sectors shall be transferred before issuing an LBA Read Continue command. The device may clear R/B# to zero prior to returning additional data. Figure 3 defines the behavior and timings for continuing the read transfer to the next Sector Multiple sectors.

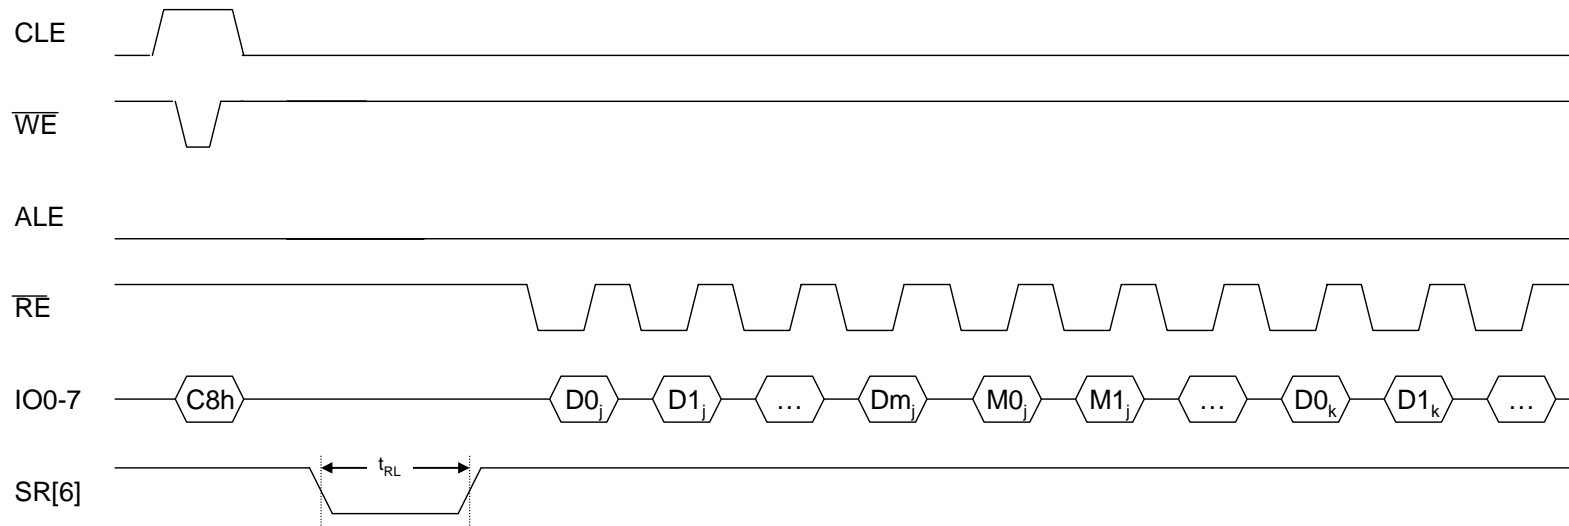


Figure 3 LBA Read Continue timing

- Dn_j Data bytes read from first sector in the Sector Multiple set of sectors.
- Dn_k Data bytes read from second sector in the Sector Multiple set of sectors.
- Mn_j Metadata bytes read for the first sector in the Sector Multiple set of sectors.

1.9. LBA Write Definition

The LBA Write function writes sectors of data identified by an LBA and a sector count to the device. If metadata is used, metadata for a sector shall be written atomically with the corresponding data sector. If metadata use is enabled, then the metadata bytes for an addressed sector are transferred immediately following the data for the sector. Thus, metadata if used is interspersed between the data sectors transferred. Figure 4 defines the LBA Write behavior and timings for the transfer of the first data block.

Read Status shall be used to check the status of LBA Write (Continue). The host shall ensure that FAIL is cleared to zero in the Status register which indicates the data write was successful.

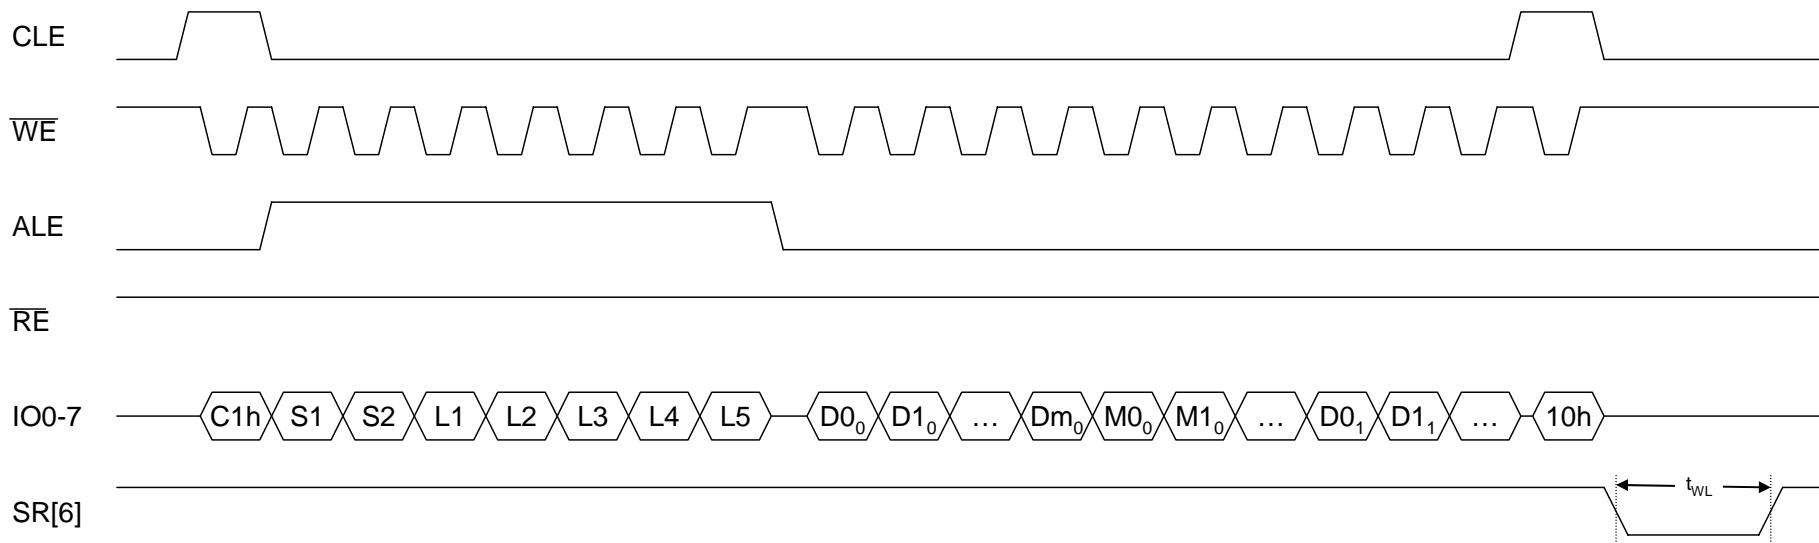


Figure 4 LBA Write timing

L1-L5 LBA of the sector(s) to write. L1 is the least significant byte.

S1-S2 Number of sectors to be written to the device. S1 is the least significant byte.

Dn₀ Data bytes written to the first addressed sector.

Dn₁ Data bytes written to the second addressed sector.

Mn₀ Metadata bytes written for the first addressed sector.

When many sectors are being transferred, the host shall issue the LBA Write Continue (C2h) command every time Sector Multiple sectors have been transferred, after R/B# has been set to one. Figure 5 defines the behavior and timings for continuing the write transfer to the next Sector Multiple sectors.

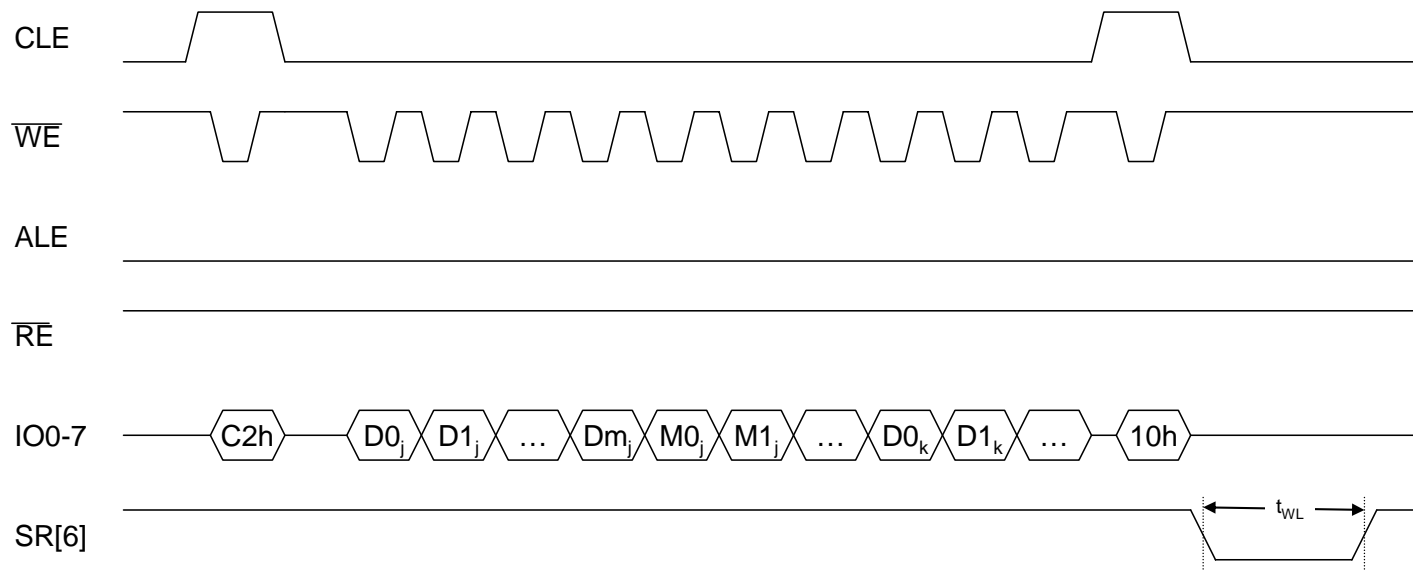


Figure 5 LBA Write Continue timing

Dn_j Data bytes written for the first sector in the Sector Multiple set of sectors.

Dn_k Data bytes written for the second sector in the Sector Multiple set of sectors.

Mn_j Metadata bytes written for the first sector in the Sector Multiple set of sectors.

1.10. LBA Read Metadata Definition

The LBA Read Metadata function returns only the metadata for the sectors identified by an LBA and a sector count. The maximum sector count specified shall be the Sector Multiple value. Metadata shall be enabled prior to use of this command. Figure 6 defines the LBA Read Metadata behavior and timings.

Read Status shall be used to check the status of the LBA Read Metadata command. The host shall ensure that FAIL is cleared to zero in the Status register which indicates the metadata read was successful prior to transfer. After completion of the Read Status command, the C0h command shall be issued by the host to continue with the data output flow.




Figure 6 LBA Read Metadata timing

L1-L5 LBA of the sector(s) to retrieve metadata for. L1 is the least significant byte.

S1-S2 Number of sectors to return metadata for. S1 is the least significant byte.

Mn₀ Metadata bytes read for the first addressed sector.

Mn₁ Metadata bytes read for the second addressed sector.

1.11. LBA Deallocate Definition

The LBA Deallocate function is issued by the host to indicate to the device an LBA range that is no longer being used by the host. This command is an optimization that allows the device to do more effective garbage collection and potentially avoids undue delays for future LBA Write commands. After issuing an LBA Deallocate for a location, the data contained in that LBA range may be destroyed by the device. Note that the device is not required to overwrite these locations with any particular value. Figure 7 defines the LBA Deallocate behavior and timings.

It is recommended that LBA Deallocate be acknowledged by the device quickly, and that the device then perform any optimizations as a background operation.

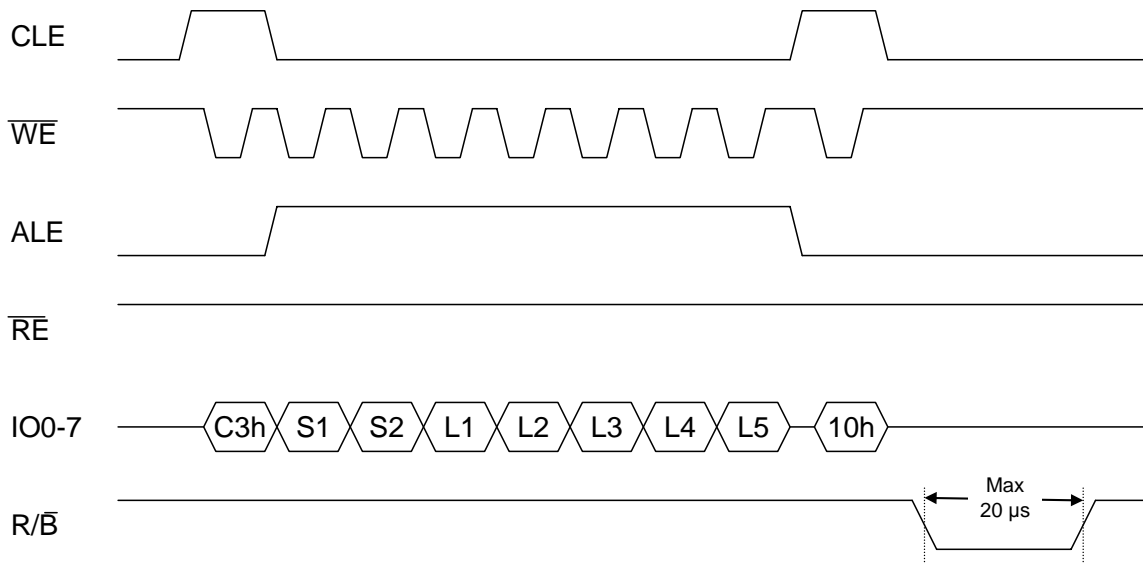


Figure 7 LBA Deallocate timing

L1-L5 LBA of the sector(s) to deallocate. L1 is the least significant byte.

S1-S2 Number of sectors to deallocate. S1 is the least significant byte.

1.12. LBA Flush Definition

The LBA Flush function ensures that all data previously written to the device is placed into the Flash array. The command shall not be completed by the device until all data is placed into the Flash array. Figure 8 defines the LBA Flush behavior and timings.

LBA Flush with the Standby parameter set to one shall be issued by the host prior to removing power from the Block Abstracted device, regardless of whether any LBA Write commands have been issued.




Figure 8 LBA Flush timing

- P1 Bit field of parameters for LBA Flush
 Bit 0: If set to one, Standby is requested after the operation completes.
 Bit 7:1 are reserved.

1.13. LBA Abort Definition

The LBA Abort command may be used to abort an outstanding LBA command, including while R/B# is cleared to zero. Specifically, LBA Abort may be used to terminate execution of an LBA Read (Continue), LBA Write (Continue), LBA Read Metadata, LBA Deallocate, or LBA Flush command. This command shall not be used to abort other commands like Get Features, Read Parameter Page, etc.

If the host does not want to transfer all sectors and/or metadata requested in a particular command sequence, the host shall issue an LBA Abort to terminate the data transfer.

When LBA Abort is used to terminate execution of LBA Read Continue or LBA Write Continue, any previously completed Sector Multiple transfer shall remain valid if success status was reported by the Block Abstracted device.

Following an LBA Abort, the status value returned by Read Status shall have FAIL set to one if a command was aborted. If FAIL is cleared to zero, then the previous command completed successfully prior to the LBA Abort being issued.




Figure 9 LBA Abort timing

1.14. Read Status Definition

The Read Status function retrieves the status value for the last operation issued. Figure 10 defines the Read Status behavior and timings.




Figure 10 Read Status timing

ST Status value as defined in section 1.16.

1.15. Read Unique ID Definition

The Read Unique ID function is used to retrieve the 16 byte unique ID (UID) for the device. The unique ID when combined with the device manufacturer shall be unique.

The UID data may be stored within the Flash array of one of the attached Flash components and directly passed through to the host. Thus, the data returned is not guaranteed to be without error. The UID is returned with its complement, as shown in Table 3. If the XOR of the UID and its bit-wise complement is all ones, then the UID is valid.

| Bytes | Value |
|-------|---------------------------|
| 0-15 | UID |
| 16-31 | UID complement (bit-wise) |

Table 3 UID and Complement

Sixteen copies of the UID and the corresponding complement shall be stored by the device. For example, reading bytes 32-63 returns to the host another copy of the UID and its complement.

Figure 11 defines the Read Unique ID behavior. The host may use any timing mode supported by the device in order to retrieve the UID data.




Figure 11 Read Unique ID command timing

$U0_k-U1_k$ The k th copy of the UID and its complement. Sixteen copies are stored. Reading beyond 512 bytes returns indeterminate values.

1.16. Status Field Definition

The returned status byte value (ST) for Read Status has the format described below. If the RDY bit is cleared to zero, all other bits in the status byte (except bit 7) are invalid and shall be ignored by the host.

| Value | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-------|----|-----|---|-----|---|-----|---|------|
| ST | na | RDY | R | VSP | R | PFR | R | FAIL |

- FAIL** If set to one, then the last command failed. If cleared to zero, then the last command was successful. Once FAIL is asserted to one for a LBA Read or LBA Write command, it shall remain set to one throughout the execution of that command. FAIL is cleared to zero when RDY is cleared to zero for an operation other than LBA Read Continue or LBA Write Continue.
- PFR** If set to one, then the device is performing a power fail recovery operation. This bit shall only be set to one following power on, and shall be cleared to zero within one minute after R/B# is initially set to one by the device.
- VSP** Vendor specific.
- RDY** If set to one, then the device is ready to accept another command and all other bits in the ST value are valid. If cleared to zero, then the last command issued is not yet complete and bits 5:0 in the ST value are invalid and shall be ignored by the host. For writes, this bit provides no indication of whether the data is written to the Flash array. The R/B# signal is equal to the value of this bit.
- R** Reserved (0)
- na** The value is undefined.

1.17. Feature Parameter Definitions

Feature settings are volatile across power cycles. For each feature setting, whether the value across resets is retained is explicitly stated.

| Feature Address | Description |
|-----------------|----------------------------|
| 00h | Reserved |
| 01h | Timing Mode |
| 02h-5Fh | Reserved |
| 60h | BA NAND: Error information |
| 61h | BA NAND: Configuration |
| 62h-7Fh | Reserved |
| 80h-FFh | Vendor specific |

1.17.1. BA NAND: Error Information

This setting shall be supported if the device supports block abstracted mode. This information shall be retained across Reset commands.

| Sub Feature Parameter | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-----------------------|--------------------|---|---|---|---|---|---|---|
| P1 | Error Information | | | | | | | |
| P2 | Health Information | | | | | | | |
| P3 | Reserved (0) | | | | | | | |
| P4 | Reserved (0) | | | | | | | |

Error Information

Bit 0: If set to one, the previous operation did not complete due to an ECC failure.

Bit 1: If set to one, the previous operation did not complete because no spare blocks were available.

Bit 2: If set to one, the host should retry the last command that failed.

Bit 7:3: Reserved.

Health Information This field indicates information regarding the current health of the implementation.

| 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|----------|---|---|---|---------------|---|-----------|---------|
| Reserved | | | | Device Status | | Read Only | Replace |

Replace When set to one, indicates that the device should be replaced and is near end of life. When cleared to zero, the device is operating normally.

Read Only When set to one, indicates that the device has reverted to read-only and can no longer accept writes. When cleared to zero, indicates that the device is read/write capable.

Device Status This field indicates the device's operational status.
 00b: Excellent
 01b: Good
 10b: Degraded
 11b: Poor

Reserved Reserved values shall be cleared to zero by the host. The device shall not be sensitive to the value of reserved fields.

1.17.2. BA NAND: Configuration

This setting shall be supported if the device supports block abstracted mode. This information is not retained across Reset commands.

| Sub Feature Parameter | 7 | 6 | 5 | 4 | 3 | 2 | 1 | 0 |
|-----------------------|--------------|---|---|---|---|---|---|-----|
| P1 | Reserved (0) | | | | | | | MDE |
| P2 | Reserved (0) | | | | | | | |
| P3 | Reserved (0) | | | | | | | |
| P4 | Reserved (0) | | | | | | | |

Metadata Enable (MDE) The metadata enable field specifies whether the host is using metadata for LBA Read and LBA Write operations. If set to one, then the host is using metadata. If cleared to zero, then the host is not using metadata. By default after power-on, MDE is zero.